



Europäisches
Patentamt

European
Patent Office

Office européen
des brevets

J1040 U.S. PRO
10/058320
01/30/02

Bescheinigung Certificate Attestation

Die angehefteten Unterlagen stimmen mit der ursprünglich eingereichten Fassung der auf dem nächsten Blatt bezeichneten europäischen Patentanmeldung überein.

The attached documents are exact copies of the European patent application described on the following page, as originally filed.

Les documents fixés à cette attestation sont conformes à la version initialement déposée de la demande de brevet européen spécifiée à la page suivante.

Patentanmeldung Nr. Patent application No. Demande de brevet n°

01102117.7

CERTIFIED COPY OF
PRIORITY DOCUMENT

Der Präsident des Europäischen Patentamts;
Im Auftrag

For the President of the European Patent Office
Le Président de l'Office européen des brevets
p.o.

R C van Dijk

DEN HAAG, DEN
THE HAGUE, 22/11/01
LA HAYE, LE



Europäisches
Patentamt

European
Patent Office

Office européen
des brevets

Blatt 2 der Bescheinigung
Sheet 2 of the certificate
Page 2 de l'attestation

Anmeldung Nr.:
Application no.: **01102117.7**
Demande n°:

Anmeldetag:
Date of filing: **31/01/01**
Date de dépôt:

Anmelder:
Applicant(s):
Demandeur(s):
GRETAG IMAGING Trading AG
5430 Wettingen
SWITZERLAND

Bezeichnung der Erfindung:
Title of the invention:
Titre de l'invention:
Automatic colour defect correction

In Anspruch genommene Priorität(en) / Priority(ies) claimed / Priorité(s) revendiquée(s)

Staat:	Tag:	Aktenzeichen:
State:	Date:	File no.
Pays:	Date:	Numéro de dépôt:

Internationale Patentklassifikation:
International Patent classification:
Classification internationale des brevets:
H04N1/62

Am Anmeldetag benannte Vertragstaaten:
Contracting states designated at date of filing: AT/BE/CH/CY/DE/DK/ES/FI/FR/GB/GR/IE/IT/LI/LU/MC/NL/PT/SE/TR
Etats contractants désignés lors du dépôt:

Bemerkungen:
Remarks:
Remarques:

EPO - Munich

69

31. Jan. 2001

Attorney's File: 53 279 VIII**GRETAG IMAGING Trading AG****Automatic Colour Defect Correction**

The present invention relates to a method for automatically correcting colour defective areas in an image. Such defective colour areas are recorded with a colour spectrum deviating from the actual colour spectrum of said areas without such colour defect. In particular, the present invention relates to an automatic red eye correction. The method of the present invention is set forth in claim 1. An image recording device or image printing device according to the invention as well as a data carrier, carrying a program representing the method of the invention, and an electromagnetic signal to control a computer device, an image printing device or an image recording device, including control data representing the method of the present invention, are referred to in the remaining independent claims.

In the field of photography and photo-finishing, the automatic correction of colour defective areas, and in particular of red eye defects, has been a subject for discussion since photographs have been taken, in particular if the subject was photographed by using artificial light, e.g. a flash light. As one example of outstanding importance, the correction of red eye defects will be discussed in this patent application, while also other colour defects can be corrected in accordance with the present disclosure.

It is known that red eye defects are caused by the reflection of flash light from the retina, which causes the pupil to look red or at least reddish. Usually, only the pupil of such a flash light photograph looks red. However, in extreme cases, also the whole eye can become a red patch and even the colour might look pinkish or a kind of orange. It is clear that such kinds of

The following references are referred to in order to incorporate the disclosure of these references with respect to the detection of the location of a face and with respect to the detection of a location of eyes. All these kinds of image processing can be used in accordance with the present invention. According to one prior art reference, the Hough transform was used for the detection of eye centers. In "Robust Eye Centre Extraction Using the Hough Transform", by David E. Benn et al, Proc. First International Conference AVBPA; pp. 3 – 9; Crans-Montana; 1997, a gradient decomposed Hough transform was used which considerably reduces the need for memory space and processing speed of a processing system for image data.

According to another approach, the flow field characteristics, which are generated by the transitions from the dark iris to the rather light sclera, were used to detect eyes in an image. As disclosed in "Detection of Eye Locations in Unconstrained Visual Images", by Ravi Kothari et al, Proc. Int. Conf. on Image Processing, ICIP 96; pp. 519-522; Lausanne; 1996, it was proposed to use these flow field characteristics. This attempt uses a field similar to an optical flow field generated for a motion analysis. Afterwards, a two-dimensional accumulator is used to achieve votes for intersections of prominent local gradients.

According to Alan L. Yuil et al, it was proposed in "Feature Extraction from Faces Using Deformable Templates", International Journal of Computer Vision, 8:2, pp. 99 – 111; 1992, to use a deformable template, which provides a model of a human eye. By minimising the costs of such a kind of fit of the template over a number of energy fields, the best fit was iteratively found.

Another kind of automatic eye detection will be referred to below, which was invented by A. Held and has been filed first with the European Patent Office under the title "Automatic Image Pattern Detection".

According to the method of the present invention, the image data stemming from the identification of the basic areas, in particular human eyes, are processed further to identify borderlines and/or centres of the colour defective areas, in particular defective red eyes. Afterwards, it is detected whether the localised area or areas, deemed to be defective, are defective or not, and in particular whether a detected eye has a red eye defect or not. If a localised area, e.g. an

In the case that it is the aim of the present invention to correct a red eye defect, and the position and size of the iris have been estimated, the maximum of the red eye defect is determined to be the actual position of the iris of an eye.

Processing to discover whether a detected human eye actually has a red eye defect, is accomplished by means of the HSV colour spaces. The abbreviations HSV stand for Hue, Saturation and Value colour spaces. Hue represents a value which corresponds to the appearance of a colour. Saturation corresponds to the amount of grey content, i.e. the amount of white or black which is existent in the colour which is to be evaluated in accordance with the HSV colour spaces. For instance, a saturation with the value zero indicates that there is no hue, but only grey scale. The value component of the HSV colour space is a measure of its brightness. The HSV colour space is normalised. Usually, a colour in the HSV colour space can be specified by stating a hue angle, a saturation level and a value level. A hue angle of zero is often defined to represent red. The hue angle increases in a counter clockwise direction and complementary colours are 180° apart. The HSV colour spaces can be device dependent or device independent.

According to this HSV colour space, a large variety of colour defective images, and in particular of images with defective red eyes, have been analysed to be projected into the HSV colour space. On the basis of such analysis, fuzzy membership functions for H, for S and for V were calculated and are used in accordance with the present invention to evaluate whether pixels or so called seed pixels are defective or not. In other words, neighbouring pixels, with respect to the actual position of the iris, are analysed considering said curves or fuzzy membership functions achieved by analysing a variety of real-world photographs with red eye defects, to acquire the fuzzy membership functions for the colour channels to be taken into account. The intersections between the three colour channel positions of each of the neighbouring pixels or seed pixels and the fuzzy membership functions are then determined and, on the basis of the intersections, it is decided whether a particular neighbouring or seed pixel is defective or not. This decision is also based on a predetermined threshold and, if the corresponding value of the pixel does not exceed the predetermined threshold, the red color content corresponds to the usual content of red colour in the image data of a usual human eye, which then does not underlie the red eye defect.

move small holes and intrusions as well as small outgrows of the correction mask. In addition, the correction mask data can be directed to binary dilation and also be influenced by a Gaussian smoothing or the like.

The, thus obtained, correction mask is then applied to the matching area of the uncorrected image to remove the red-eye defect or any other colour defect which has to be corrected. At this stage of the method, the correction mask is represented by a grey-scale mask which allows for comparatively efficient gradual corrections towards the borderlines of the colour defective area, and in particular the red-eye defective area. The corresponding correction which can be achieved gives an appearance of the corrected image which is close to the natural appearance of the corrected motive, in particular a human eye.

Since, according to the invention, it is preferred that the correction mask provide normalised grey scale data in a range from $m=0\dots1$, it is possible to use the following Equation to correct the uncorrected input image:

$$R_{\text{new}} = R \cdot m(R - \min(G, B)) \quad (\text{Eq. 1.3})$$

If, for instance, there is no defect at all for a pixel, then the corresponding value m in the correction mask, i.e. the correction factor m , is 0 as well. Otherwise, according to Equation 1.3, the red colour channel of this particular pixel will be decreased towards the minimum of the green colour channel and the blue colour channel.

According to another embodiment of the invention, if the green colour channel or the blue colour channel of a particular pixel are discovered to be very large, i.e. the difference between the green colour channel and the blue colour channel is really large, the bigger of the two colour channels will have to be adjusted as well to avoid another colour shift which could render the appearance of the defective area unnatural after the correction. This further necessary adjustment can be implemented by Equation 1.3 as well.

An image processing device for processing image data, which can implement the method according to the present invention, includes an image data input section, an image data processing section and an image data recording section for recording processed image data. Such

estimate calculated to be within the borders of the initial estimate.

Fig. 8 shows fuzzy membership functions calculated with respect to the HSV colour space.

Fig. 9a shows a correction mask before a smoothing operation.

Fig. 9b shows the correction mask according to Fig. 9a after a smoothing operation.

Fig. 10 is a flow diagram showing the principles of the method according to the present invention.

Fig. 11 shows Sobel operators to be used in an embodiment of the invention.

Fig. 12 is a flow diagram depicting a first stage of the method in accordance with one embodiment of the invention.

Fig. 13 shows a pictogram of a face.

Fig. 14 shows a pictogram of a human eye.

Fig. 15 shows one embodiment of a second stage of an embodiment of the method of the present invention.

Fig. 16 shows the distribution as a result of one embodiment of the first stage of the invention.

Fig. 17 shows the distribution according to Fig. 16 after further processing.

While the method according to the invention can also be used to correct colour defects of many other kinds, it will be discussed here with reference to the automatic detection and correction of red eye defects in the field of photo finishing. A red-eye defect occurs if a very strong light, e.g. a flash light is reflected from the retina, i.e. the back of the eye. The red colour is indicative for the blood vessels in the retina. In terms of colorimetry, red-eye defects are not defined properly and there is a very thin borderline between defective and not defective eyes. In general, in red-defective eyes, however, the values for the red colour channel will be considerably higher than the values for the two other colour channels (green, blue).

As shown in Fig. 1, there are several modules that, in turn, reduce a domain to be searched for. The domain to be searched, for instance, is a basic area in an uncorrected image, in which there is an increased likelihood of finding a particular colour defective area, and in terms of this description, a red eye defect. There can be two detection stages, automatic face detection and automatic eye detection, which are general-purpose modules. This means, these modules

tion of faces happens fully automatic and that the detection rate is reasonably high and the false negative rate, that is, faces being detected even though there is no face present, is reasonably low. What reasonable constitutes will depend on the actual context of the application. The disclosure of the Rowley and the Schneiderman references is incorporated into this application.

As most face detectors are not invariant to rotation, it can be useful to ensure that all the possible orientations of faces can be detected. How to do this will highly depend on the face detector being used, as the rotation invariance of each detector will vary widely. For instance, in Rowley's approach, rotation invariance is given within approximately $\pm 15^\circ$. On the other hand, in the approach by Schneiderman, rotation invariance is given in a range of about $\pm 45^\circ$. Therefore, rotation invariance has to be ensured by external means, this can for instance be done by pre-rotation of the image, followed by a post-processing and the normal face detection. This is shown in Fig. 3.

For a system based on the face detector by Schneiderman, four stages are necessary. In other words, the face detector is applied to images rotated by 0° , 90° , 180° and 270° , respectively.

Once a face has been detected, the search space for finding eyes can be restricted considerably. According to the above-described method, it is possible to obtain a bounding box of a face, together with its approximate orientation. As stated before, face detectors are, in general, not rotation invariant. Therefore, orientation of the face could be obtained in the range given by the rotational invariance of the face detector, which could be up to $\pm 45^\circ$ in the case of the Schneiderman detector.

There is a variety of approaches that can be applied for detecting eyes. Again, similar to the detection of faces, it is important to have an approach that works fully automatic, has a high recognition rate, and a low false positive rate.

The basic approach for automatic eye detection is outlined in Fig. 4. As pre-processing step, any processing can be incorporated that will enhance facial features, as for instance, histogram normalisation, local contrast enhancement, or even red-enhancement according to Equation (1.1) for red-eye detection. In general, it is a good idea to normalise the input image,

Fig. 10 shows a flow diagram for the automatic detection of image patterns and particularly for human eyes, the sun, a flashlight reflection or the like. The detection is carried out in two stages: a coarse stage followed by a refinement stage. During the coarse stage, the exact locations of the searched image pattern are of less interest. However, attention is rather directed to areas that are of interest and that are likely to contain the searched image patterns, e.g. eyes. During the refinement stage those regions will then be further examined and it will then be determined whether there actually is a searched image pattern, e.g. an eye and, if yes, what is its location and approximate size.

In the following, the disclosure is directed to the recognition of the location of eyes, while it is, of course, possible to proceed with other image patterns approximately the same way.

For both the coarse and the refinement detection stage, the gradient decomposed Hough transform is relied on for the detection of eyes.

The classical theory of the Hough transform will be referred to below. This transform is the classical method for finding lines in raster images. Consider the equation of a line in Equation (2.1).

$$y = mx + c \quad (2.1)$$

If, for each set pixel in the image, x and y are kept fixed and a line is drawn in the accumulator space according to Equation (2.2), then for each line that is formed in the original image, all the lines drawn in the accumulator will intersect in one place, namely the place that determines the proper parameters for that line in question.

$$c = xm + y \quad (2.2)$$

The original theory of the Hough transform can be extended to accommodate other curves as well. For instance, for circles, it is possible to use the parameter model for a circle as given in Equation (2.3). Now, however, this will require a three-dimensional parameter space.

First, it is reasonable to start the approach for the detection of eyes with some kind of pre-processing. Here, for instance, it is useful to normalise the input image to a known size, given by a model face image, or any kind of histogram normalisation or local contrast enhancement can be performed. For this approach described here, it's preferred to restrict the domain of the input by only looking at a part of the image. Assuming that the input image is a proper face image, preferably the output from some face detection scheme, it is decided to look only at the upper 2/3 of the image as shown in Fig. 4. This will allow to neglect parts of the mouth and even the nose, that contain a lot of curved features and could mislead further detection of the eyes.

Depending on the domain of the system, which is further processed, it is useful to apply some special colour space conversions in order to stress certain features. For instance, if eyes for later red-eye removal are to be detected, it is useful to employ a red-enhanced colour space as input to the gradient calculations, as is shown in Equation (3.1).

$$I_{red} = R - \min(G, B) \quad (3.1)$$

Given the pre-processed input image, it is possible to proceed to calculate the gradient information, which will then be needed for the actual Hough transform. The gradient images can either be calculated by applying Sobel templates or operators as shown in Fig. 11, or by utilising other gradient information, as for instance can be obtained from the Canny edge detector.

At this stage, it is decided to apply a straight-line removal procedure to the gradient images. This will allow the influence of very strong, but straight, gradients on the accumulator to be reduced considerably. The outline of straight-line removal is shown in Fig. 13. Straight-line removal attempts to isolate straight lines from the detected edges and removes those areas from the gradient image. In general, this will result in a much better detection of the eye center.

Straight-line removal as shown in Fig. 12, includes the following steps. First, the edges of the image are extracted by applying some edge detector, for instance, the Canny edge detector.

The input to the second stage, i.e. the refinement stage, are the isolated boxes or surroundings from the previous stage, each containing a possible eye candidate, together with the gradient images as described before. An outline of the refinement stage is given in Fig. 15.

Basically, the approach is the same as for the coarse detection stage. However, instead of having one two-dimensional accumulator, now two one-dimensional accumulators are used. This means, each accumulator will contain the projection of all the votes onto the axis in question. Differently to the coarse detection stage, where a projection would incur many spurious peaks due to spatial ambiguities, in the case of the eye boxes, it can safely be assumed that there is not more than one object of interest within the surrounding or box. Therefore, using projections will considerably simplify the task of actually fitting a model to the accumulator, as it has only to deal with one-dimensional functions. Again, the projections would look somewhat similar to the cross-section as shown in Figs. 16 and 17, and they can be treated accordingly, following Equation (3.2). For the remaining values in the accumulator, a Gaussian distribution can be used and its mean and standard deviation can be calculated. The two means, one from the x projection and one from the y projection, directly give the location of the eye center. The minimum of the two standard deviations will be taken as an estimate for the size of the eye.

For the projection onto the x-axis, the estimate of location and size will be rather accurate in general, due to the symmetry. For the projection onto the y-axis, however, there might be some kind of bias if there is a strong eyebrow present. In practice, however, the influence of this can be neglected, as it usually will be offset by other gradient edges below the eye.

For each detected eye candidate, it is possible to further extract some kind of confidence measure by looking at how many votes this position received in the two-dimensional accumulator space. A high number of votes strongly corroborates the actual presence of an eye.

According to the invention, an automatic approach to image pattern detection based on the hierarchical application of a gradient decomposed Hough transform has been presented. Due to the splitting up of the task into a coarse and a fine stage, it is possible to get a much more

After the strongest defect (New Estimate) has been detected, this strongest defect is used as the center for the further extension to neighbouring pixels directly located around the center or center pixel. The neighbouring pixels or seed pixels are then deemed to represent a layer of neighbouring or seed pixels.

Whether a seed pixel really belongs to a red-eye defect is determined according to its HSV value and a fuzzy membership function as shown in Fig. 8. Actual values for the parameters are given in Table 1. If the fuzzy intersection r of the three channels has been calculated, see Equation (1.2), then it can be decided with a simple thresholding operation whether the seed pixel belongs to the class of red-eye defects or not

$$r = \frac{\text{hsv}}{\max(h,s,v)} \quad (\text{Eq. 1.2})$$

Based on the so selected seed pixels, one starts to recursively add layers of further neighbouring or seed pixels around the respective last layer. First, all the background pixels that border the seed region are marked. For each such pixel, a test is done as to whether its HSV values permit it to be added to the seed region. For this test, the same set of fuzzy membership functions as before can be used, as shown in Fig. 8.

The model from Fig. 8 was obtained by analysing a number of red-eye defects from various real-world photographs. Again, by taking the fuzzy intersection of the HSV values, it can be decided by thresholding whether or not to add the pixel in question to the seed region, which, after its completion will provide the correct mask. In addition, if the pixel in question is an edge pixel as indicated by the edge image, this pixel is omitted. After adding one layer to the seed region, it is tested whether its major geometric properties are still within the specified limits. Basically, the created seed region should not exceed a certain size and its eccentricity should be closed to a circle. The latter can be tested by taking the ratio between the major and minor axes of the seed region. If, during this iteration, some neighbouring pixels were added and the region passed the geometry test then we proceed to adding a new layer. Otherwise, the region growing stops.

According to the present disclosure, a hierarchical approach has been used that consecutively reduces the search space by taking into account semantical information about the image in question. In addition, the hierarchical approach allows one already at an early stage to abandon the processing of images which do not contain the necessary context. Such a system can be operated in a fully automatic manner, allowing for incorporation into a completely unsupervised system.

Attorney's File: 53 279 VIIIEPO - Munich
69

31. Jan. 2001

Claims

1. Method for automatically correcting colour defective areas in an image, said defective colour areas being recorded with a colour spectrum deviating from the actual colour spectrum of said areas without colour defects, in particular automatic red eye correction, wherein:
 - a) basic areas in the image are identified on the basis of features which are common for these recorded defective areas, e.g. skin detection to identify skin related areas, face detection on the basis of a face pictogram or the like and/or an eye detection, and so on, said basic areas supporting an increased likelihood of including defective areas, in particular red eye defects;
 - b) the processing is reduced to the basic areas to identify borderlines and/or centers of the defective areas, in particular red eye defects;
 - c) it is identified whether the localised area or areas, deemed to be defective, is/are defective or not, in particular whether a detected eye has a red eye defect or not;
 - d) if a localised area, e.g. an eye, has been identified to be defective, a correction mask is created to correct the visual appearance of the defective area.
2. Method according to claim 1, characterized in that, alternatively to the feature 1a), a specialised portrait mode can be activated in the case of defective red eye detection.
3. Method according to claim 1, wherein if, after one of the steps 1a) to 1c), the processing of an image has not resulted in the identification of an area to be processed further, the processing of an image is stopped and another image can be processed.
4. Method according to one of claims 1 to 3, wherein after a center of a basic area, e.g. an eye, has been identified, an expected defect is emphasised in its intensity (I , I_{red}), to more easily detect whether the basic area is defective, wherein in the case of a detected colour defect, in particular a defective red eye, the following Equation is particularly useful:

$$I_{red} = R - \min(G, B)$$

(Eq. 1.1),

9. Method according to claim 8, wherein the analysis of, and extension to, neighbouring pixels is terminated if no further other neighbouring pixels have been identified as red-eye defective and/or if the borderlines of the defective red eye have been reached or exceeded.
10. Method according to one of claims 8 or 9, wherein the pixel data representing the correction mask are directed to at least one smoothing operation.
11. Method according to one of claims 8 to 10, wherein the correction mask, preferably a grey-scale correction mask, is applied to the matching area of the uncorrected image to remove the colour defect, and in particular the red-eye defect.
12. Method according to claim 11, wherein the removal of the colour defect, and in particular the red-eye defect, is practised by means of the following Equation:

$$L_{red\ new} = R - m(R - \min(G, B)) \quad (\text{Eq. 1.3}),$$

where m is an element of the correction mask corresponding to a particular pixel of the uncorrected image.

13. Method according to one of claims 11 or 12, wherein, if a particular pixel of the uncorrected image has a considerably large difference between the green channel and the blue channel, the larger channel is adjusted in accordance with Equation (1.3).
14. Image processing device for processing image data, including:
 - a) an image data input section,
 - b) an image data processing section,
 - c) an image data recording section for recording processed image data,
 - d) wherein the image data processing section implements a method according to one of claims 1 to 13.
15. Computer data signal embodied in carrier wave, comprising instructions to implement the method according to one of claims 1 to 13.

Attorney's File: 53 279 VIIIEPO - Munich
69
31. Jan. 2001**Abstract**

The present invention refers to a method for automatically correcting colour defective areas in an image, which defective colour areas were recorded with a colour spectrum deviating from the actual colour spectrum of said areas without colour defects, wherein basic areas in the image are identified on the basis of features which are common for these recorded defective areas, said basic areas supporting an increased likelihood to include defective areas, and the processing is then reduced to the basic areas to identify borderlines and/or centres of the defective areas, and afterwards, it is identified whether the localised basic area or areas deemed to be defective are defective or not, and finally, if a localised basic area has been identified to be defective, a correction mask is created to correct the visual appearance of the defective area.

EPO - Munich

69

31. Jan. 2001

Fig. 1

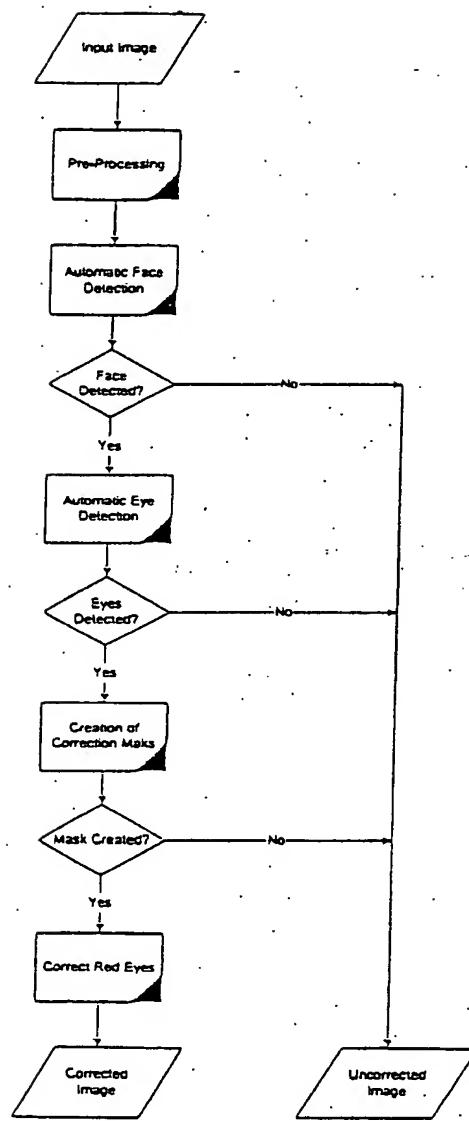


Fig. 3

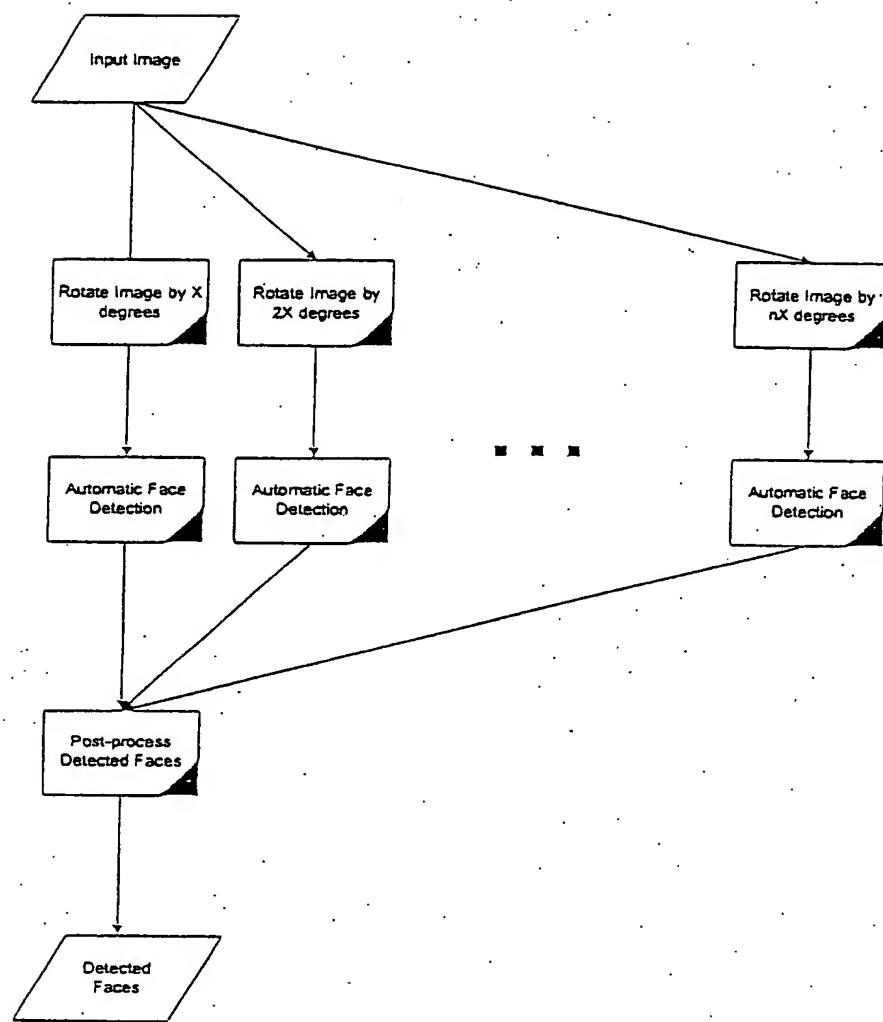


Fig. 6

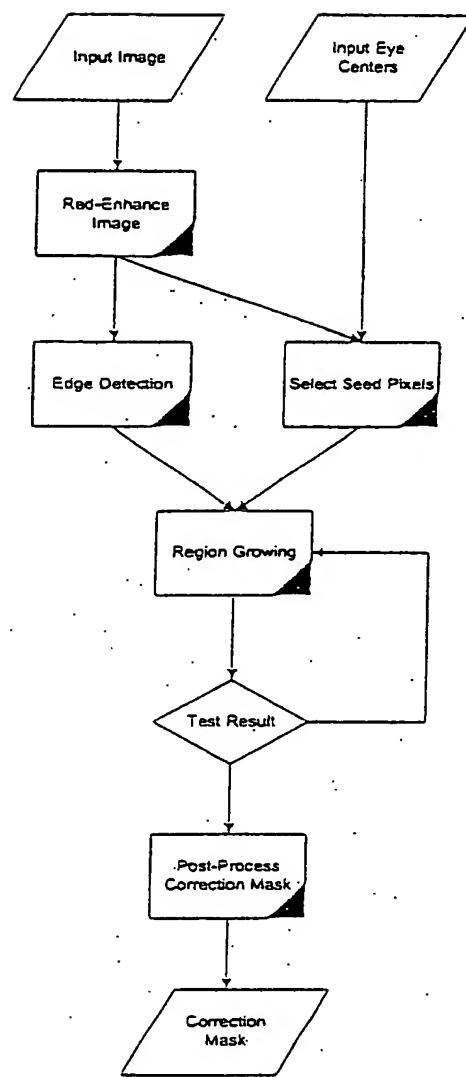
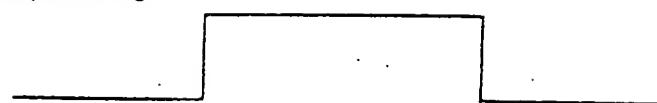


Fig. 9a

Before Smoothing

**Fig. 9b**

After Smoothing

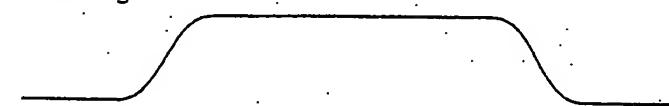


Fig. 11

$$\begin{matrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{matrix} \quad \begin{matrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{matrix}$$

Fig. 12

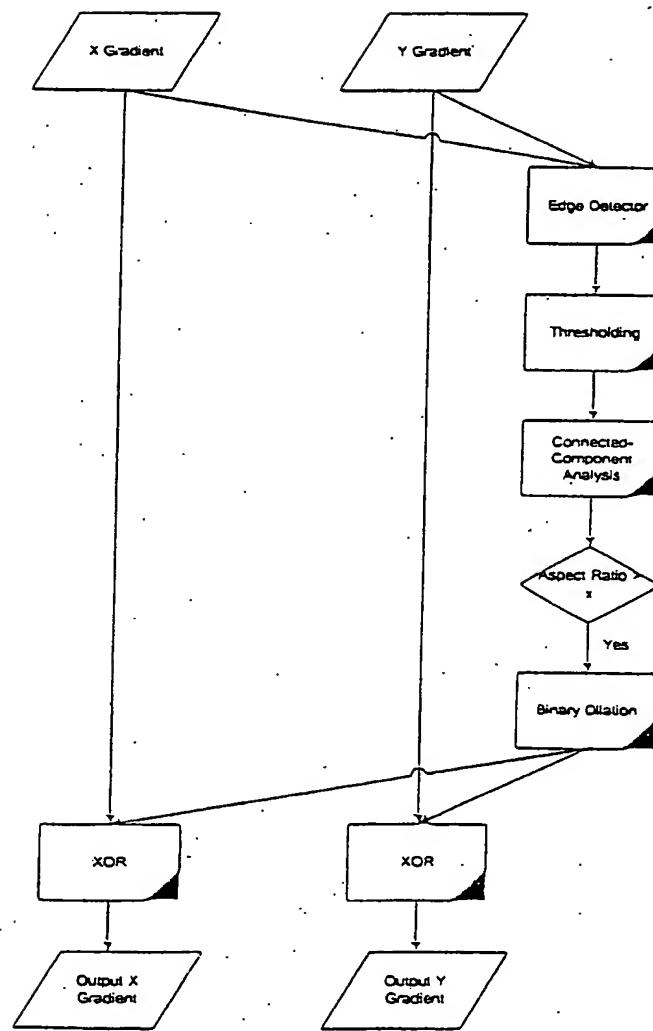


Fig. 15

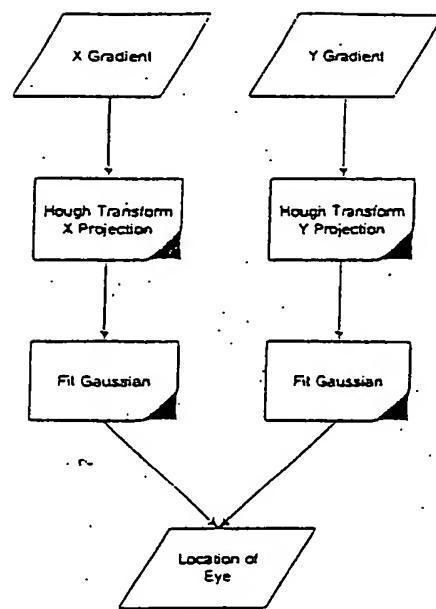


Fig. 16

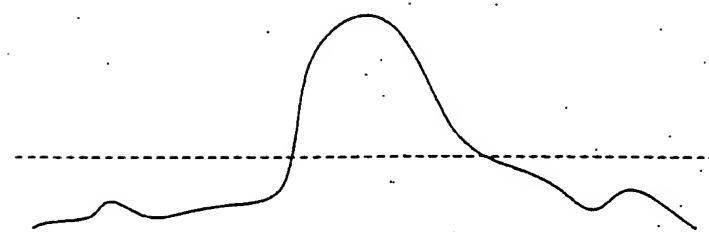


Fig. 17

